LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-PROC-416866

Nuclear Forces and High-Performance Computing: The Perfect Match

T. Luu, A. Walker-Loud

September 16, 2009

SciDAC 2009
San Diego, CA, United States
June 15, 2009 through June 18, 2009

Disclaimer

Nuclear Forces and High-Performance Computing: The Perfect Match

T Luu[1], **A Walker-Loud**[2]

[1] N-division, Lawrence Livermore National Laboratory, Livermore, CA 94551 USA

[2] Physics Department, College of William and Mary, Williamsburg, VA 23187-8795 USA

E-mail: `tluu@llnl.gov`

Abstract. High-performance computing is now enabling the calculation of certain nuclear interaction parameters directly from Quantum Chromodynamics, the quantum field theory that governs the behavior of quarks and gluons and is ultimately responsible for the nuclear strong force. We briefly describe the state of the field and describe how progress in this field will impact the greater nuclear physics community. We give estimates of computational requirements needed to obtain certain milestones and describe the scientific and computational challenges of this field.

1. Introduction

The nuclear force acts between particles called hadrons. Hadrons exist as either quark/anti-quark pairs called mesons, such as pions and kaons, or triplets of quarks called baryons, such as neutrons and protons. The range of the force is very small, acting over distances less than 2 fm (1 fm=10^{-15} m). Its diminutive range is contrasted by its enormous strength. For example, at these same distances the attractive part of the force readily binds neutrons and protons to form nuclei, despite the electromagnetic repulsion between protons. Further, at distances $<$ .5 fm, the attractive part of the force quickly gives way to an even stronger repulsive core. Figure 1 gives a schematic representation of the nuclear potential between neutrons and protons.

Aside from forming the nuclei that compose the periodic table of elements, the nuclear force plays an integral role in the evolution of our universe. For example, just minutes after the big bang, the nuclear force was responsible for the synthesis of primordial elements lighter than ^{7}Lithium. These elements spawned the first stars that utilized the nuclear force to burn light-ion elements by fusion. From the ashes of this burning came heavier elements, ranging from the life-giving carbon element to elements as heavy as iron and nickel. In the death throes of massive stars undergoing supernova, the repulsive core of the force is ultimately responsible for the final core bounce which generates the shock wave that ultimately ejects the stars' mantel. In this process, the synthesis of even heavier elements are made, such as the actinides, through the 'r-process', the rapid capture of neutrons on nuclei. In more mundane settings, the nuclear force, in a delicate balance with the electromagnetic interaction, is responsible for the process of fission where nuclei split into smaller constituents, a process exploited in nuclear power stations to produce carbon-free energy. ***In all subfields of nuclear physics the nuclear force plays a central role.***

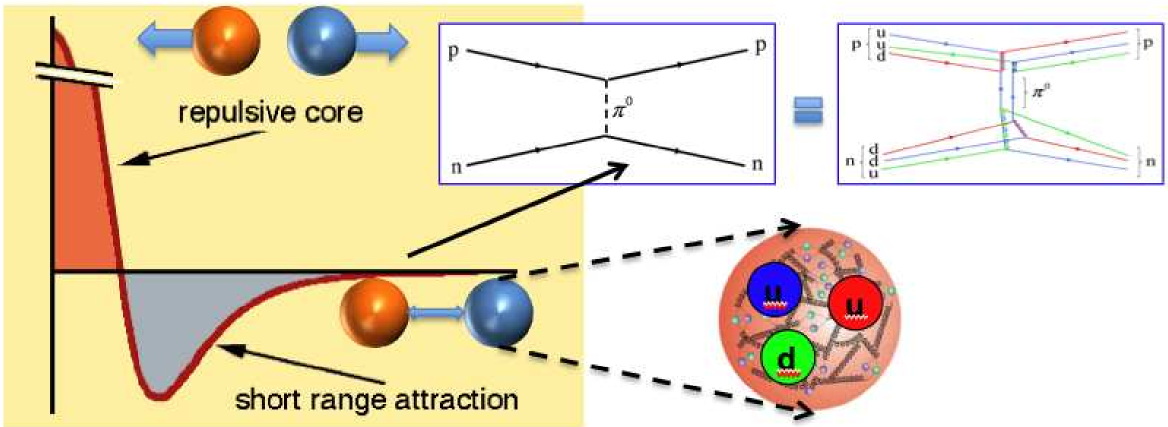

Figure 1. *Schematic diagram of the two-body potential between neutrons and protons. The neutrons and protons themselves are composite particles consisting of quarks and gluons. At the fundamental level, the force between neutrons and protons is a complicated process involving the interactions between quarks and gluons.*

1.1. Open questions about the nuclear force

There are many aspects of the nuclear force, however, that are poorly understood even today. This is due, in part, to the fact that many systems of interest in nuclear physics are not experimentally accessible. At the same time, nuclear scientists have historically lacked the theoretical and computational sophistication to calculate properties of nuclear systems, such as their force, from first principles. Below is just a sample of the types of questions plaguing this field:

- ***What is the force between hyperons?***
  We do not know the nature of the nuclear force between hyperons, exotic hadrons containing strange quarks. *Is the force attractive? Repulsive?* Understanding this force, even at this level, will impact nuclear astrophysics since these particles can be present in, for example, the cores of neutron stars. The nature of their interactions will impact the equation of state governing the properties of neutron stars, and therefore its evolution.
- ***Why is the nuclear spin-orbit force so large?***
  Nuclei exhibit shell structure akin to that of atomic electron orbitals of elements that form the periodic table. Nuclear scientists attribute this shell structure, and the 'magic numbers' associated with it, to a nuclear spin-orbit force. However, the strength of this force (determined empirically) is much larger than conventional wisdom. Determining the reasoning behind this will help us better understand the collective behavior of neutrons and protons that form nuclei.
- ***Where does the tensor force come from?***
  The deuteron represents the simplest nuclear system, consisting of a bound proton and neutron. It has a magnetic dipole moment and electric quadrapole moment, which suggests the presence of a nuclear tensor force. Even with today's modern nuclear many-body algorithms, the calculated moments still differ from experimental measurements. A better understanding of the origin of this force will not only improve our calculations, but will gain us deeper knowledge of nature's intricate fine tunings at the nuclear level.

- ***What is the strength and origin of the parity-violating nuclear force?***
  The abundance of matter over anti-matter in our universe remains one of the largest unsolved puzzles in modern physics. The known parity-violating forces of the standard model are not sufficient to explain the observed baryon asymmetry of the universe. Searching for beyond-the-standard model parity violating physics first requires a precise knowledge of the parity violating interactions predicted by the standard model. Lattice QCD can be used to compute parity violating matrix elements arising both explicitly from the standard model as well as from proposed higher dimensional operators.
- ***What is the nature of the three-nucleon force?***
  In contrast to the electromagnetic force, the nuclear force can have a pure three-body interaction. There is strong theoretical and empirical evidence for the existence of this interaction, but to date we have not been able to determine its exact nature. The lack of knowledge of this interaction represents the largest source of uncertainty in nuclear structure and reaction calculations. A complete understanding, for example, of the life-enabling synthesis of carbon that occurs in stars (*i.e.* the Hoyle process) can only be accomplished with a better understanding of the three-nucleon force. Determining this force will impact broader fields of physics, let alone the nuclear physics community.

1.2. Quantum Chromodynamics

Our *inability* to answer these questions stems from the fact that the nuclear force is not a fundamental force; rather, it is a manisfestation of an underlying force between quarks and gluons described by a quantum field theory known as quantum chromodynamics (QCD). This theory governs the interactions of quarks and gluons that are basic constituents that make up hadrons, such as pions, neutrons, and protons. At high energies ($E \gg 1$ `GeV`$= 10^9$ `eV`), QCD is represented by a coupled set of non-linear equations that can be linearized. This perturbative nature has allowed for direct comparison with high energy experiments, giving scientists insight into nature's workings over distances that are smaller than the size of neutrons and protons ($\ll 1$ `fm`$= 10^{-15}$ `m`). However, at low energies or larger distances, the theory becomes strongly coupled, which has so far precluded any analytic solutions. Thus, the theoretical determination of fundamental nuclear physics phenomena directly from QCD has been met with less success. A long-standing effort of the U.S. Department of Energys (DOE) Nuclear Physics program is understanding how QCD in this low-energy regime manifests itself into observed nuclear phenomena.

The strongly coupled aspect of QCD equates to the phenomenon of confinement: quarks, through their interaction with gluons, are never found in isolation at these low energies; they combine to form hadrons. The phenomenon of confinement is a well established outcome of QCD[1], but the exact nature of how confinement occurs and all of its consequences are still unresolved. This mystery is what precludes us from answering many fundamental questions in nuclear physics. In principle, if we can solve QCD in this energy regime we will be understand this mystery, and therefore answer these (and many more) questions.

1.3. Lattice QCD

Lattice QCD (LQCD) is a formulation of QCD in which space and time are discretized on a lattice and the theory is confined to a periodic box. In this manner the theory becomes amenable to numerical calculations. It is directly applicable to the low-energy regime in which QCD is analytically insoluble. A lattice QCD calculation is performed using a Monte Carlo method in which samplings of the QCD vacuum are generated with a distribution prescribed by

[1] The converse of confinement, *asymptotic freedom*, where quarks become non-interacting at distances $\ll$1 `fm`, was the subject of the 2004 Nobel Price in Physics awarded to D. Gross, H. Politzer, and F. Wilczek.

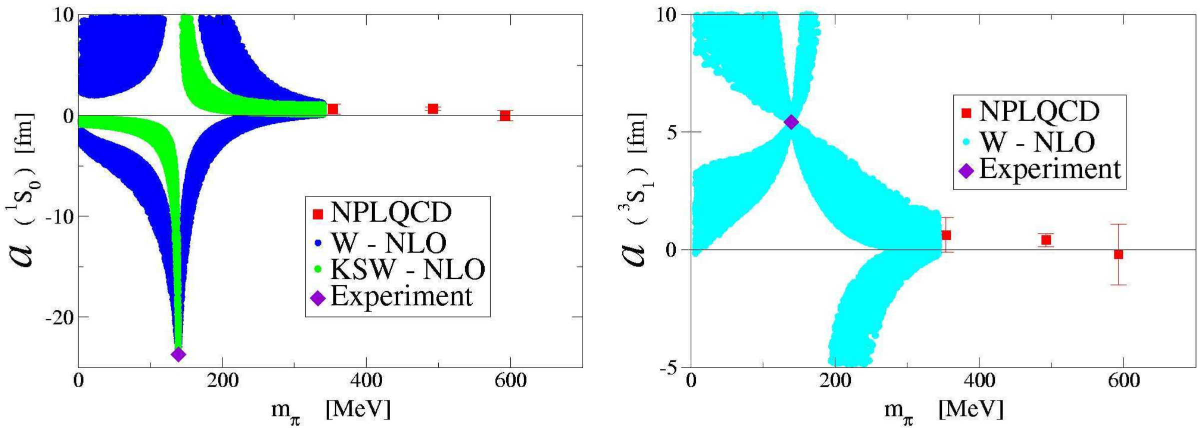

Figure 2. *Calculation of NN scattering lengths [3].*

QCD, and physical observables are then measured on these samplings. The greater the number of measurements, the smaller the statistical uncertainty in the calculation. With sufficient computational resources, and a careful account of the discretization artifacts imposed by the lattice, calculations that are not currently possible with pencil and paper are now numerically feasible. To date, LQCD offers the only means of rigorously calculating QCD observables in the low-energy regime relevant to nuclear physics. It represents the best opportunity for answering longstanding questions in nuclear physics. Lattice QCD, however, is highly computational intensive.

2. The impact of high-performance computing on nuclear forces so far

The US Lattice QCD collaboration (USQCD) [1] represents the largest contingent of lattice gauge practitioners within the US. It supports US scientists who develop and use large-scale, high-performance computers for calculations in LQCD through SciDAC grants. With support from the DOE, it designs, constructs and operates large-scale computing systems for LQCD calculations. Allocations on these large-scale computing systems are granted through a peer-reviewed process in coordination with the DOE's INCITE program [2].

The recent large strides made in LQCD can be attributed in large part to the cohesive program brought about by USQCD and its coupling with SciDAC. The examples below, in the area of nuclear forces, were made possible by allocations from USQCD and the INCITE program. These examples are by no means exhaustive.

- ***The nuclear force is perturbative at large*** m_π
  Lattice QCD is a tool to calculate physical observables such as particle masses and excited states, or in the case of nuclear forces, the scattering lengths a, effective ranges r, and in principle the full scattering phase shift $\delta(p)$. The first dynamical calculation of nucleon-nucleon (NN) scattering with lattice QCD was only recently carried out at three values of the pion mass with $m_\pi \geq 350$ MeV [3]. It was found that at these pion masses, the interaction was weakly repulsive in both the $^1S_0$ and $^3S_1$ channels, see Fig. 2. This is in contrast to the same system at the physical pion mass, where we know that, for example, the deuteron is (weakly) bound and has an unnaturally large scattering length a compared to its effective range r, *i.e.* $a/r \gg 1$. *Lattice QCD calculations tell us how nuclear phenomena depend on the fundamental parameters of nature.*
- ***Weinberg's prediction of pion interactions works well for all values of*** m_π
  The simplest system of interacting hadrons, both theoretically and numerically, is that of

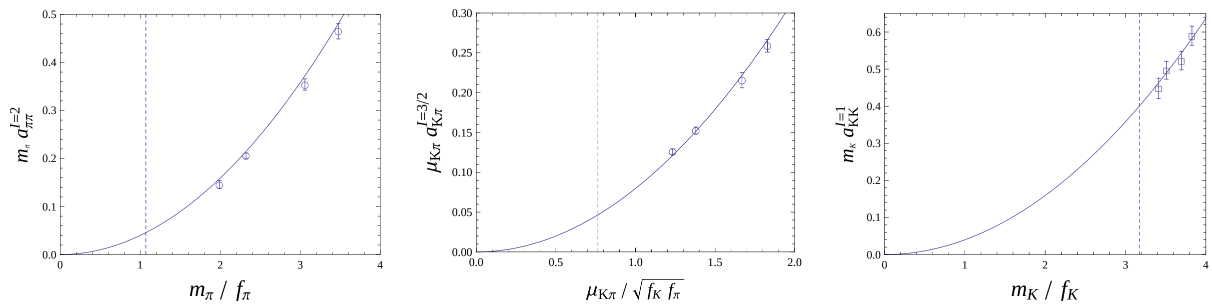

(a) (b) (c)

Figure 3. *Calculated values of the $I = 2$ $\pi\pi$ (a), $I = 3/2$ $K\pi$ (b) and $I = 1$ KK (c) scattering lengths. In all cases, the numerical-data are plotted vs. the lattice-physcal masses and decay constants, as calculated for each ensemble. The dashed lines represent the physical point and the curves are Weinberg's LO prediction. Note the apparent absence of quantum loop corrections, which would presumably shift the lattice results away from the LO prediction.*

two mesons, such as two pions or two kaons. In 1966 Weinberg predicted the leading order contribution to their interaction [4]. Recent LQCD calculations of these systems, which include all contributions to the force and not just Weinberg's leading order term, found that even at larger-than-physical m_π Weinberg's prediction works surprisingly well (there is no significant evidence of extra 'quantum loop' corrections in these processes), see Fig.3. Extrapolation to the $m_\pi \to 0$ 'chiral' limit was performed [5, 6, 7] to predict the strength of the force at the infinite volume, continuum, physical pion mass limit (*i.e.* the 'physical point'), finding in the case of two charged pions (at the charged pion mass) [8, 9],

$$m_\pi a_{\pi\pi}^{I=2} = 0.04330 \pm 0.00042 \,. \tag{1}$$

Similar calculations have been performed on systems comprised of a charged pion and kaon [10] and two charged kaons [11]. *In certain cases, LQCD calculations can obtain a precision at the sub-percent level.*

- ***We now know that pions exhibit a pure, repulsive three-body force***
  The most powerful aspect of LQCD is its predictive capability. With sufficient computer resources, systems not accessible by experiment can be calculated on a lattice. In most cases, these systems have impact in many areas of nuclear physics, such as nuclear astrophysics. A recent example of this predictive capability came from LQCD simulations of multiple pions a box. A careful analysis of these systems allowed Detmold et al. to extract a three-body force between pions, finding that it is repulsive with a size consistent with naive dimensional analysis [12, 13]. A subsequent work determined the equivalent three-body interaction for charged kaons [14] finding an interaction strength consistent with zero. This year, we have seen the first lattice calculations comprising three baryon systems with lattice QCD [15]. *For the first time, LQCD is allowing definitive statements of hadronic three-body forces to be made.*

3. Why HPC and nuclear forces make a perfect match

Lattice QCD calculations of nuclear physics observables, such as nuclear interaction parameters, involve multiple algorithms that utilize varying degrees of computational resources. Here we loosely outline, from top to bottom, the steps taken in a LQCD calculation:

- ***Generation of lattice gauge ensembles***
  As mentioned in sect. 1.3, samples of the QCD vacuum, called lattice gauge configurations, must be generated. The generation of these ensembles involve high-dimensional integrals that are performed by hybrid Monte Carlo techniques. Space-conserving molecular dynamics integrators, such as the Omelyan integrator [16], are used to update Monte Carlo timesteps and rejection/acceptance is done by importance sampling. At every step of the molecular dynamics loop, large matrix inversions must be performed. These inversions represent the most computational demanding aspect of the calculation, and are typically done by conjugate-gradient methods coupled recently with deflation techniques [17]. The dimension of these matrices, as well as their condition number, typcially range in the 10^7-10^8 [18], and only get worse as the pion mass is lowered. The resources needed to generate these ensembles roughly scale as [19]
  $$b^{-6}L^5m_\pi^{-3}\ ,$$
  where b is the lattice spacing and L is the length of the side of a box. Leadership-class machines are most suited for these types of calculations.
- ***Measurements performed on these lattices***
  With the ensembles generated, measurements must be taken on the lattices that comprise the ensembles. These measurements typically take the form of propagators and involve large matrix inversions as well. Conjugate gradient techniques, and their variants, coupled with deflation techniques are used. Leadership-class machines are most suited for these calculations.
- ***Contraction of propagators***
  The propagators must be contracted to form physically relevant objects. This procedure is combinatoric, scaling roughly as $n_u!n_d!n_s!$, where n_u represents the number of up quarks, n_d the number of down quarks, and n_s the number of strange quarks. In the past this procedure has been done in serial, but as the number of baryons investigated increases, the factoral growth will preclude serial calculations.
- ***Data analysis***
  The large amounts of data generated, coupled with the stochastic nature of the data, means that robust data analysis techniques will be needed to extract statistically meaningful quantities, such as arguments of exponents. Recently a new technique has been proposed to extract such data that utilizes the Prony method [20, 21]. In the future, other methods need to be developed (such as Bayesian analysis) to allow for verification and validation. These calculations can be done on clusters.

The methods described above show how, in the process of performing a LQCD calculation relevant for nuclear forces, a substantial fraction of problems associated HPC is involved. Further, the stochastic, determistic, and combinatoric issues apply to various hardware platforms ranging from dedicated leadership-class machines to large-scale clusters to simple serial machines. From a computer science and applied mathematics perspective, investing research into nuclear physics and LQCD offers the ideal platform for addressing many issues in HPC, while at the same time impacting a fundamental area of physics.

4. Scientific and computational challenges

To fully take advantage of HPC in the coming years, certain scientific and compuatational challenges need to be addressed in the area of nuclear forces. Though none of these issues below represent 'show-stoppers' to progress by themselves, they do present serious obstacles that can impede the current rate of advancement of this field. There is active research in developing algorithms to address many of these issues.

- ***Signal-to-noise ratio***
  Current LQCD calculations are inherently stochastic. Calculations of baryon systems, such as the deuteron, suffer from poor signal-to-noise ratios. This impedes any extraction of multi-baryon interaction parameters. As computational resources increase, signal-to-noise issues will diminish slightly, but only with development of novel algorithms and computational techniques can the signal-to-noise issue be laid to rest. Recently a procedure was proposed to ameliorate this issue through modification of the box boundary conditions [22, 23].
- ***Scaling multi-baryon codes for high-performance capability***
  Existing algorithms for performing multi-baryon calculations are not suited for large scale computing, and will have to be revamped to take advantage of high-performance computing capability as we approach the exascale era. Novel algorithms for multi-baryon systems need to be developed to fully utilize large computational resources.
- ***Development of finite-volume effective field theories***
  Continued development of finite volume effective field theories (EFTs) need to take place so that Lattice QCD calculations can be matched onto theories utilized by other areas of nuclear physics, such as nuclear structure and reactions, and nuclear astrophysics. Such theories will must include full pion dynamics, endowing them with full quark-mass dependence on scattering parameters. This will allow for the most robust extrapolation methods.
- ***Interfacing with large scale nuclear structure calculations***
  Lattice QCD calculations of few-nucleon interaction parameters will ultimately be fed into nuclear many-body calculations [24] via the use of EFTs. This will entail substantial collaboration with the nuclear structure community–something that currently does not exist. Both theoretical and numerical methods need to be developed to enhance the overlap between Lattice QCD and the nuclear structure and reactions community, as well as other areas of physics.
- ***Disconnected diagrams***
  To date, Lattice QCD calculations have generally involved a certain class of calculations, those of connected diagrams, due to computational limitations. These diagrams represent the propagation of quarks from the initial to the final states. A complete description of parity violation in nuclear structure will include forces that arise at short distances which involve disconnected diagrams.
- ***Requirements needed to perform lattice measurements***
  Lattice QCD measurements of nuclear physics observables on configuration ensembles require commensurate computational resources as those used to generate the ensembles themselves. Further, these measurements require larger memory allocations. As we approach the exascale era, these requirements are currently estimated to be at least two orders of magnitude greater than todays available resources.

5. Nuclear force milestones and their requisite resources

A recent DOE-sponsored nuclear physics workshop on extreme-scale computing [25] culminated in a timeline for milestones in nuclear physics as a function of computational resources. We list here the milestones relevant to nuclear forces, as shown in Fig. 4.

- ***Precision meson-meson interactions***
  Already precise calculations of certain interaction parameters of various meson systems are being performed, *e.g.* see sect. 2. Other aspects of the force between mesons will be probed with HPC resources of at least .1 petaflop-year sustained. As the calculation of disconnected

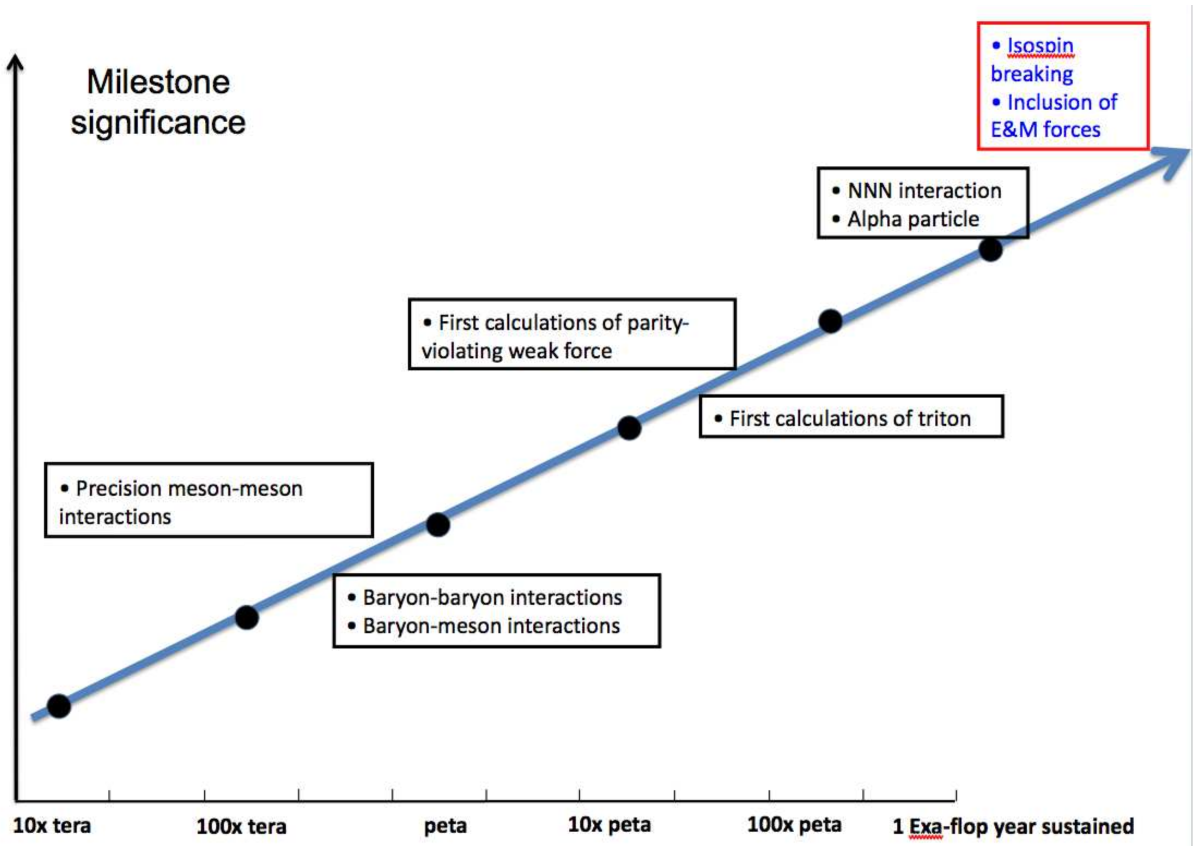

Figure 4. *Anticipated milestones in the area of nuclear forces and the required computational resources.*

diagrams becomes available, more types of meson systems can be probed.
`Resources needed:` $>$`.1 petaflop-year sustained.`

- ***Precision baryon-baryon and baryon-meson interactions***
  Calculations of baryon interaction parameters suffer from poorer signal-to-noise ratios than their mesonic counterparts. Significantly more measurements must be performed to counter address this issue. Further, the large combinatoric scaling involved with contractions compounds the issue further.
  `Required resources:` $>$ `1 petaflop-year sustained.`
- ***Parity-violating interaction term*** $h_{\pi NN}$
  As the calculation of disconnected diagrams becomes available, as well as the inclusion of class of algorithms that involve four-point functions, the first components of the nuclear parity-violating force, $h_{\pi NN}$ will be calculated.
  `Required resources:` $>$ `10 petaflop-year sustained.`
- ***The triton system***
  An essential step in extracting the nuclear three-body force will be the initial calculation of the triton, a nuclear system composed of a single proton and two neutrons.
  `Required resources:` $>$ `100 petaflop-year sustained.`
- ***Precision nuclear three-body force***
  Only after a full mapping of the properties of the triton system, including its excited states in a box, can a precise extraction of the three-body force be performed.

`Required resources: >.1 to 1 exaflop-year sustained.`

- ***Isospin-breaking forces and inclusion of the electromagnetic force***
  Lattice QCD calculations are currently performed at the 'isospin limit' where the masses of the up and down quark are equal. Further, electromagnetic forces are ignored. In reality, the masses of the up and down quarks are not equal, and since the quarks themselves have electromagnetic charges they are subjected to electromagnetic forces. Including these effects in LQCD calculations will require exascale resources and beyond.
  `Required resources: > 1 exaflop-year sustained.`

These rough estimates are made using current experience and knowledge about lattice calculations. It is highly probable that disruptive technologies and algorithms will accelerate the timescale for achieving these goals.

6. The integration of high-performance computing with nuclear forces

In the future the full integration of high-performance computing with nuclear forces will benefit the nuclear physics community in the following manner:

- ***Predictive capability will become the norm in nuclear physics***
  Historically nuclear physics has been empirically driven. High-performance computing, coupled with LQCD calculations, will endow nuclear physics with predictive prowess of unprecedented scope. Systems not accessible by experiment will be calculated from first principles with precision that constantly improves in lock-step with gains in HPC.
- ***The various subfields of nuclear physics will operate with increased synergism***
  The increased computational resources of the past decade have brought the various subfields of nuclear physics to a point where inter-disciplinary research is becoming a reality. No longer is it possible for different subfields of nuclear physics to exist independently from another, as the level of precision of calculations has reached a stage where continued improvement involves direct collaboration with other areas of physics. The import of HPC on each individual subfield of nuclear physics is clear. Just as important will be the holistic impact of HPC on the entire nuclear physics community. Figure 5 shows the various connections that will be strenghthened due to HPC from the viewpoint of nuclear forces.
- ***Lattice QCD will become one of the most powerful calculational tools available to nuclear scientists***
  Because QCD represents the fundamental theory of the nuclear force, future HPC calculations of nuclear physics observables using LQCD will set the standard for studies not only in theoretical nuclear physics, but also experimental nuclear physics.
- ***The next generation of scientists will be experts in both HPC and nuclear physics***
  The integration of HPC with the field of nuclear forces and the broader nuclear physics community offers an ideal platform in which to train the next generation of scientists that are both experts in HPC and nuclear physics. Already, through efforts supported by SciDAC, we are witnessing the coupling of seemingly disparate fields, such as computational science, applied mathematics, and nuclear physics, to solve grand challenges in physics. This coupling will get stronger, and must become stronger to fully utilize HPC as one approaches the exascale. Succesful scientists in this era will be the ones that can work seamlessly bewteen these different fields.

7. Conclusion

High-performance computing is bringing nuclear physics into a new era where predictive capability will become the norm. In this process, answers to longstanding questions in

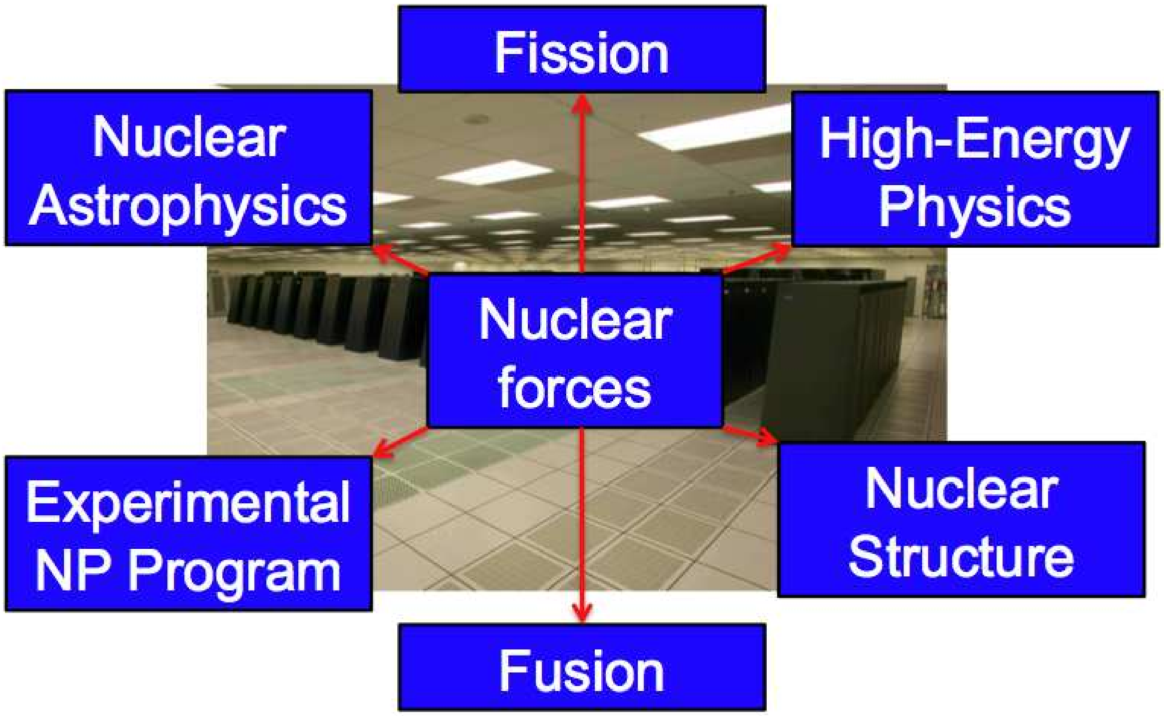

Figure 5. *Examples of some cross-cutting threads related to nuclear forces due to the integration of high-performance computing.*

fundamental physics will be attainable, and the connection between different subfields of nuclear physics will become stronger. Nowhere is this more realizable than in the area of nuclear forces, where high-performance computing is now allowing aspects of hadronic interactions to be ascertained from true *ab intio*, first principles LQCD calculations. The impact of these calculations will be far reaching, affecting the fields of nuclear structure and nuclear astrophysics, for example. In this sense, nuclear forces and high-performance computing represent the perfect match.

In these proceedings we described recent progress in the area of nuclear forces with respect to HPC, and give estimates on computational resources needed to obtain important milestones. We also outlined the benefits of HPC to the nuclear physics community as a whole. As we argued in these proceedings, LQCD promises to be calculational tool for nuclear physics of unprecedented power and scope. High-performance computing will make this promise come true.

Acknowledgments
We thank members of NPLQCD for their inciteful critiques of this manuscript. Luu thanks M. Hruška for useful discussions. Walker-Loud was supported under OJI Grant DE-FG02-07ER41527. Luu performed this work under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

References

[1] USQCD: URL `http://www.usqcd.org/`
[2] INCITE: URL `http://www.er.doe.gov/ascr/incite/`
[3] Beane S R, Bedaque P F, Orginos K and Savage M J 2006 *Phys. Rev. Lett.* **97** 012001 (*Preprint* `hep-lat/0602010`)

[4] Weinberg S 1966 *Phys. Rev. Lett.* **17** 616–621
[5] Chen J W, O'Connell D, Van de Water R S and Walker-Loud A 2006 *Phys. Rev.* **D73** 074510 (*Preprint* hep-lat/0510024)
[6] Bedaque P F, Sato I and Walker-Loud A 2006 *Phys. Rev.* **D73** 074501 (*Preprint* hep-lat/0601033)
[7] Chen J W, O'Connell D and Walker-Loud A 2007 *Phys. Rev.* **D75** 054501 (*Preprint* hep-lat/0611003)
[8] Beane S R, Bedaque P F, Orginos K and Savage M J (NPLQCD) 2006 *Phys. Rev.* **D73** 054503 (*Preprint* hep-lat/0506013)
[9] Beane S R *et al.* 2008 *Phys. Rev.* **D77** 014505 (*Preprint* 0706.3026)
[10] Beane S R *et al.* 2006 *Phys. Rev.* **D74** 114503 (*Preprint* hep-lat/0607036)
[11] Beane S R *et al.* (NPLQCD) 2008 *Phys. Rev.* **D77** 094507 (*Preprint* 0709.1169)
[12] Beane S R *et al.* 2008 *Phys. Rev. Lett.* **100** 082004 (*Preprint* 0710.1827)
[13] Detmold W *et al.* 2008 *Phys. Rev.* **D78** 014507 (*Preprint* 0803.2728)
[14] Detmold W, Orginos K, Savage M J and Walker-Loud A 2008 *Phys. Rev.* **D78** 054514 (*Preprint* 0807.1856)
[15] Beane S R *et al.* 2009 (*Preprint* 0905.0466)
[16] Omelyan I P, Mryglod I M and Folk R 2003 *Comput. Phys. Comm.* **151** 272–314
[17] Luscher M 2004 *Comput. Phys. Commun.* **156** 209–220 (*Preprint* hep-lat/0310048)
[18] Joo B these proceedings
[19] Christ N H, Dumitrescu T T, Loktik O and Izubuchi T 2007 *PoS* **LAT2007** 351 (*Preprint* 0710.5283)
[20] Fleming G T, Cohen S D, Lin H W and Pereyra V 2009 (*Preprint* 0903.2314)
[21] Beane S R *et al.* 2009 *Phys. Rev.* **DXX** (*Preprint* 0903.2990)
[22] Bedaque P F and Walker-Loud A 2008 *Phys. Lett.* **B660** 369–375 (*Preprint* 0708.0207)
[23] Bedaque P F and Walker-Loud A 2008 (*Preprint* 0811.2127)
[24] Vary J these proceedings
[25] DOE Extreme-scale workshop: URL http://extremecomputing.labworks.org/nuclearphysics/index.stm